

(19)



Europäisches Patentamt

European Patent Office

Office européen des brevets



(11)

EP 0 757 333 A2

(12)

## EUROPEAN PATENT APPLICATION

(43) Date of publication:  
05.02.1997 Bulletin 1997/06

(51) Int Cl.<sup>6</sup>: G06T 9/00

(21) Application number: 96305775.7

(22) Date of filing: 05.08.1996

(84) Designated Contracting States:  
DE FR GB NL SE

• Wynn, Aaron S.  
Palo Alto, California 94301 (US)

(30) Priority: 04.08.1995 US 511326

(74) Representative: Harris, Ian Richard et al  
c/o D. Young & Co.,  
21 New Fetter Lane  
London EC4A 1DA (GB)

(71) Applicant: SUN MICROSYSTEMS, INC.  
Mountain View, CA 94043 (US)

(72) Inventors:  
• Deering, Michael F.  
Los Altos, California 94024 (US)

## (54) 3D image decoding

(57) Three-dimensional compressed geometry is decompressed with a unit having an input FIFO receiving compressed data bits and outputting to an input block state machine and an input block, whose outputs are coupled to a barrel shifter unit. Input block output also is input to Huffman tables that output to the state machine. The state machine output also is coupled to a data path controller whose output is coupled to a tag decoder, and to a normal processor receiving output from the barrel shifter unit. The decompressor unit also includes a position/color processor that receives output from the barrel shifter unit. Outputs from the normal processor and position/color processor are multiplexed to a format converter. For instructions in the data stream that generate output to the format converter, the decompression unit generates a tag sent to the tag decoder in parallel with bits for normals that are sent to the format converter. The decompressed stream of triangle data may then be passed to a traditional rendering pipeline, where it can be processed in full floating point accuracy, and thereafter displayed or otherwise used.

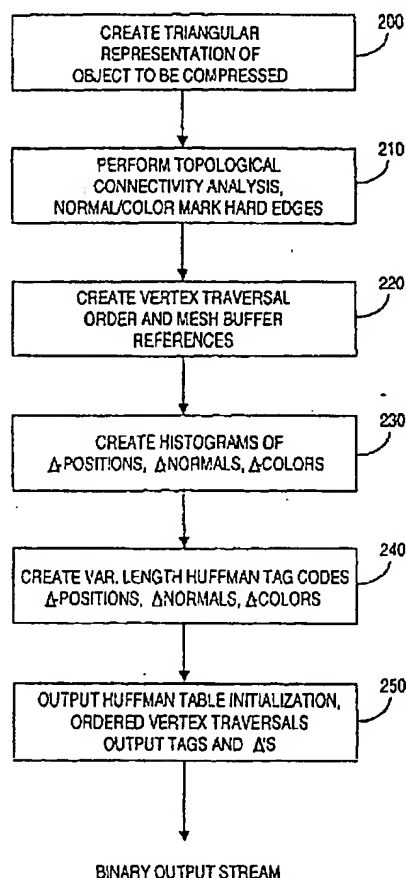


FIGURE 5

## Description

The present invention relates generally to decompressing graphics data, and more particularly to methods and apparatuses that decompress compressed three-dimensional geometry data.

Modern three-dimensional computer graphics use geometry extensively to describe three-dimensional objects, using a variety of graphical representation techniques. Computer graphics find wide use in applications ranging from computer assisted design ("CAD") programs to virtual reality video games. Complex smooth surfaces of objects can be succinctly represented by high level abstractions such as trimmed non-uniform rational splines ("NURBS"), and often detailed surface geometry can be rendered using texture maps. But adding more realism requires raw geometry, usually in the form of triangles. Position, color, and normal components of these triangles are typically represented as floating point numbers, and describing an isolated triangle can require upwards of 100 bytes of storage space.

Understandably, substantial space is necessary for three-dimensional computer graphics objects to be stored, e.g., on a computer hard disk or compact disk read-only memory ("CD-ROM"). Similarly, considerable time is necessary for such objects to be transmitted, e.g., over a network, or from disk to main memory.

Geometry compression is a general space-time trade-off, and offers advantages at every level of a memory/interconnect hierarchy. A similar systems problem exists for storage and transmission of two-dimensional pixel images. A variety of lossy and lossless compression and decompression techniques have been developed for two-dimensional pixel images, with resultant decrease in storage space and transmission time. Unfortunately, the prior art does not include compression/decompression techniques appropriate for three-dimensional geometry, beyond polygon reduction techniques. However, the Ph.D. thesis entitled *Compressing the X Graphics Protocol* by John Danskin, Princeton University, 1994 describes compression for two-dimensional geometry.

Suitable compression can greatly increase the amount of geometry that can be cached, or stored, in the fast main memory of a computer system. In distributed networked applications, compression can help make shared virtual reality ("VR") display environments feasible, by greatly reducing transmission time.

Most major machine computer aided design ("MCAD") software packages, and many animation modeling packages use constructive solid geometry ("CSG") and free-form NURBS to construct and represent geometry. Using such techniques, regions of smooth surfaces are represented to a high level with resulting trimmed polynomial surfaces. For hardware rendering, these surfaces typically are pre-tessellated in triangles using software before transmission to rendering hardware. Such software pre-tessellation is done even on hardware that supports some form of hardware NURBS rendering.

However, many advantages associated with NURBS geometric representation are for tasks other than real-time rendering. These non-rendering tasks include representation for machining, interchange, and physical analysis such as simulation of turbulence flow. Accurately representing trimming curves for NURBS is very data intensive, and as a compression technique, trimmed NURBS can not be much more compact than pre-tessellated triangles, at least at typical rendering tessellation densities. Finally, not all objects are compactly represented by NURBS. Although many mechanical objects such as automobile hoods and jet turbine blades have large, smooth areas where NURBS representations can be advantageous, many objects do not have such areas and do not lend themselves to such representation. Thus, while NURBS will have many applications in modelling objects, compressed triangles will be far more compact for many classes of application objects.

Photo-realistic batch rendering has long made extensive use of texture map techniques to compactly represent fine geometric detail. Such techniques can include color texture maps, normal bump maps, and displacement maps. Texture mapping works quite well for large objects in the far background, e.g., clouds in the sky, buildings in the distance. At closer distances, textures work best for three-dimensional objects that are mostly flat, e.g., billboards, paintings, carpets, marble walls, and the like. More recently, rendering hardware has begun to support texture mapping, and real-time rendering engines can also apply these techniques.

However, texture mapping results in a noticeable loss of quality for nearby objects that are not flat. One partial solution is the "signboard", in which a textured polygon always swivels to face the observer. But when viewed in stereo, especially head-tracked VR stereo, nearby textures are plainly perceived as flat. In these instances, even a lower detail but fully three-dimensional polygonal representation of a nearby object would be much more realistic.

Polyhedral representation of geometry has long been supported in the field of three-dimensional raster computer graphics. In such representation, arbitrary geometry is expressed and specified typically by a list of vertices, edges, and faces. As noted by J. Foley, et al. in *Computer Graphics: Principles and Practice*, 2nd ed., Addison-Wesley, 1990, such representations as winged-edge data structures were designed as much to support editing of the geometry as display. Vestiges of these representations survive today as interchange formats, e.g., Wavefront OBJ. While theoretically compact, some compaction is sacrificed for readability by using ASCII data representation in interchange files. Unfortunately, few if any of these formats can be directly passed as drawing instructions to rendering hardware.

Another historical vestige in such formats is the support of N-sided polygons, a general primitive form that early rendering hardware could accept. However, present day faster rendering hardware mandates that all polygon geometry

be reduced to triangles before being submitted to hardware. Polygons with more than three sides cannot in general be guaranteed to be either planar or convex. If quadrilaterals are accepted as rendering primitives, it is to be accepted that they will be arbitrarily split into a pair of triangles before rendering.

Modern graphics languages typically specify binary formats for the representation of collections of three-dimensional triangles, usually as arrays of vertex data structures. Thus, PHIGS PLUS, PEX, XGL, and proposed extensions to OpenGL are of this format form, and will define the storage space taken by executable geometry.

It is known in the art to isolate or chain triangles in "zigzag" or "star" strips. For example, Iris-GL, XGL, and PEX 5.2 define a form of generalized triangle strip that can switch from a zigzag to star-like vertex chaining on a vertex-by-vertex basis, but at the expense of an extra header word per vertex in XGL and PEX. A restart code allows multiple disconnected strips of triangles to be specified within one array of vertices.

In these languages, all vertex components (positions, colors, normals) may be specified by 32-bit single precision IEEE floating point numbers, or 64-bit double precision numbers. The XGL, IrisGL, and OpenGL formats also provide some 32-bit integer support. The IrisGL and OpenGL formats support vertex position component inputs as 16-bit integers, and normals and colors can be any of these as well as 8-bit components. In practice, positions, colors, and normals can be quantized to significantly fewer than 32 bits (single precision IEEE floating point) with little loss in visual quality. Such bit-shaving may be utilized in commercial three-dimensional graphics hardware, providing there is appropriate numerical analysis support.

However compressed, geometric data including three-dimensional geometry data must be decompressed to be useful. For example, applicant's patent application filed on the same date as this application, entitled METHOD AND APPARATUS FOR GEOMETRIC COMPRESSION OF THREE-DIMENSIONAL GRAPHICS DATA, discloses such compression.

Thus, there is a need for method and apparatus for decompressing three-dimensional geometry that has been compressed. Preferably, decompression is such that the output data may be passed to rendering hardware directly as drawing instructions. Finally, decompression of three-dimensional geometry should be implementable using hardware, software, or a combination thereof.

Aspects of the invention are set out in the accompanying independent claims. Preferred and further features of the invention are set out in the dependent claims. Different combinations of the features of the dependent claims may be made, as appropriate, with the features of the independent claims.

For decompression according to an embodiment of the present invention, three-dimensional geometry is first represented as a generalized triangle mesh, which allows each instance of a vertex in a linear stream to specify an average of between 1/3 triangle and 2 triangles. Individual positions, colors, and normals are quantized, with a variable length compression being applied to individual positions, colors, and normals. Quantized values are delta-compression encoded between neighbors to provide vertex traversal orders, and mesh buffer references are created. Histograms of delta-positions, delta-normals and delta-colors are created, after which variable length Huffman tag codes, as well as delta-positions, delta-normals and delta-colors are created. The compressed output binary stream includes the output Huffman table initializations, ordered vertex traversals, output tags, and the delta-positions, delta-normals, and delta-colors.

Decompression of such compressed three-dimensional geometry data may be implemented in hardware, software, or a combination of each. The decompression unit includes an input FIFO that receives compressed data bits and a signal noting size of the incoming data. The FIFO outputs are coupled to an input block state machine and an input block. Outputs from the input block and input block state machine are coupled to a barrel shifter unit. Input block output also is input to Huffman tables that output to the state machine. The state machine output also is coupled to a data path controller whose output is coupled to a tag decoder, and to a normal processor receiving output from the barrel shifter unit. The decompressor unit also includes a position/color processor that receives output from the barrel shifter unit. Outputs from the normal processor and position/color processor are multiplexed to a format converter.

For instructions in the data stream that generate output to the format converter, the decompression unit generates a 12-bit tag that is sent to the tag decoder in parallel with bits for normals that are sent to the format converter. A read-back path is used to read back the internal state of the decompressor unit. The decompressor unit carries out the following procedures:

- (1) Fetch the rest of the next instruction, and the first 8 bits of the following instruction;
- (2) Using the tag table, expand any compressed value fields to full precision;
- (3A) If values are relative, add to current value; otherwise replace;
- (3B) If mesh buffer reference, access old values;
- (3C) If other command, do housekeeping;
- (4) If normal, pass index through ROM table to obtain full values;
- (5) Output values in generalized triangle strip form to next stage.

The decompressed stream of triangle data may then be passed to a traditional rendering pipeline, where it can be processed in full floating point accuracy, and thereafter displayed or otherwise used.

Other features and advantages of the invention will appear from the following description in which the preferred embodiments have been set forth in detail, by way of example only, in conjunction with the accompanying drawings, in which:

FIGURE 1 depicts a generalized network system over which compressed three-dimensional geometry may be transmitted for decompression, according to the present invention, at the receiving end;

FIGURE 2 depicts a generalized triangular mesh data structure, and generalized mesh buffer representation of surface geometry;

FIGURE 3 depicts six-way sign-bit and eight-way octant symmetry in a unit sphere, used to provide forty-eight way reduction in table look-up size;

FIGURE 4A depicts a vertex command in a geometry compression instruction set;

FIGURE 4B depicts a normal command in a geometry compression instruction set;

FIGURE 4C depicts a color command in a geometry compression instruction set;

FIGURE 4D depicts a mesh buffer reference command in a geometry compression instruction set;

FIGURE 4E depicts a set state instruction in a geometry compression instruction set;

FIGURE 4F depicts a set table command instruction in a geometry compression instruction set;

FIGURE 4G depicts a pass through command instruction in a geometry compression instruction set;

FIGURE 4H depicts a variable length no-op command instruction in a geometry compression instruction set;

FIGURE 4I depicts tag and  $\Delta$ -position data structure;

FIGURES 4J-1 and 4J-2 depict alternative tag and  $\Delta$ -normal data structure;

FIGURE 4K depicts tag and  $\Delta$ -color data structure;

FIGURE 5 is a flowchart of method steps in a geometry compression algorithm;

FIGURE 6 is a simplified block diagram of an embodiment of decompressor hardware, according to the present invention;

FIGURES 7A-7L depict objects compressed under different conditions;

FIGURE 8 is a detailed overall block diagram of an embodiment of a decompressor unit, according to the present invention;

FIGURE 9 is a detailed block diagram of the input block shown in Figure 8;

FIGURE 10 is a detailed block diagram of the barrel shifter unit shown in Figure 8;

FIGURE 11 is a detailed block diagram of the position/color processor unit shown in Figure 8;

FIGURE 12A is a detailed block diagram of the normal processor unit shown in Figure 8;

FIGURE 12B is a detailed block diagram showing the decoder, fold, and ROM look-up components associated with the normal processor unit of Figure 12A;

FIGURE 13 is a block diagram showing interfaces to a mesh buffer, as shown in Figure 11 and/or Figure 12A;

FIGURE 14A depicts interfaces to Huffman tables, according to the present embodiment;

FIGURE 14B depicts a preferred format for entry of the Huffman table data, according to the present embodiment;

FIGURE 15A depicts a vertex instruction, according to the present embodiment;

FIGURE 15B depicts vertex component data formats, according to the present embodiment;

FIGURE 15C depicts the format for the set normal instruction, according to the present embodiment;

FIGURE 15D depicts a set color instruction, according to the present embodiment;

FIGURE 15E depicts a mesh buffer reference instruction, according to the present embodiment;

FIGURE 15F depicts a set state instruction, according to the present embodiment;

FIGURE 15G depicts a set table instruction, according to the present embodiment;

FIGURE 15H depicts a passthrough instruction, according to the present embodiment;

FIGURE 15I depicts a variable-length NOP instruction, according to the present embodiment; and

FIGURE 15J depicts a skip 8 instruction, according to the present embodiment.

A graphics decompressor according to an embodiment of the present invention decompresses three-dimensional graphics objects. Three-dimensional compression of such geometry advantageously permits a reduction in the time needed to transmit the compressed three-dimensional geometry, e.g., over a network, as well a reduction of the space wherein the geometry may be stored, e.g., on a CD-ROM, or the like.

Before describing decompression of compressed three-dimensional graphics, the overall environment in which the present invention may be practiced will be described with respect to Figure 1.

Figure 1 depicts a generalized network over which three-dimensional compressed geometry data may be transmitted, and decompressed using software, hardware, or a combination of each at the receiving end. Of course, decompression of three-dimensional graphics compression according to the present invention may be practiced upon compressed data that is presented other than via a network, e.g., compressed data stored in a memory, on a CD-ROM, and the like.

As shown in Figure 1, a source of three-dimensional graphics data 10 may be coupled to a server or encoder system 20 whose processed and compressed output is coupled over one or more networks 30 to one or more target clients or decoder systems 40. The network may be homogeneous, heterogeneous, or point-to-point.

Server 20 includes a central processing unit 50 that includes a central processor unit per se ("CPU") 60 with associated main memory 70, a mesh buffer 80, a memory portion 90 that may include a compression algorithm, and a region of read-only-memory ("ROM") 100. Alternatively, compression accordingly may be carried out in part or wholly in hardware as opposed to software. Server 20 also includes a three-dimensional graphics compression unit 60, whose compressed output data is arranged by a disk layout unit 70 for storage onto storage disk unit 80, which may include one or more CD-ROMs. The server communicates over the network(s) 30 via network interface unit 110. Those skilled in the art will appreciate that server 20 may include a mechanism for arbitrating between a plurality of client-decoder requests for compressed data.

As described in applicant's patent application filed on the same date as this application, entitled METHOD AND APPARATUS FOR GEOMETRIC COMPRESSION OF THREE-DIMENSIONAL GRAPHICS DATA, lossy compression of three-dimensional geometric data can produce ratios of 6:1 to 10:1, with little loss in displayed object quality. The following portions of this Specification will describe compression, as set forth in the above-referenced patent application, to facilitate a better understand of decompression, according to the present invention.

In a network environment, at the receiving end, decoder system(s) 40 include a network interface unit 120, a unit 130, according to the present embodiment, that decompresses three-dimensional graphics data, and whose output is coupled to a three-dimensional graphics rendering unit 140. System 40 further comprises a central processing system 150 that includes a CPU 160, memory 170, a portion of which 180 may include decompression software, and ROM 190. Compressed three-dimensional graphics may advantageously be decompressed using software, hardware, or a

combination of each. The decompressed output from decoder 40 further may be coupled to a viewer 200, or to another system requiring the decompressed graphics. Of course, unit 40 may be a standalone unit, into which precompressed three-dimensional graphics are coupled from storage 82, disks or CD-ROM 84, or the like, for decompression. Unit 40 may, for example, comprise a computer or workstation.

Assuming that three-dimensional graphics compression unit 60 functions as described in applicant's above-noted patent application, triangle data will have first been converted into a generalized triangle mesh. For a given fixed capacity of storage medium 80, a triangle mesh data structure is a near-optimal representation of triangle data. In the preferred embodiment, three-dimensional graphics object may be represented as three-dimensional triangular data, whose format after conversion causes each linear strip vertex, on average, to specify from about 1/3 triangles to about 2 triangles. Further, such triangle strip structure permits extraction of the compressed geometry by a single monotonic scan over the vertex array data structure.

Figure 2 depicts a generalized triangular mesh data structure, and generalized mesh buffer representation of surface geometry. Such a mesh data structure may be used in three-dimensional geometry compression, although by confining itself to linear strips, a generalized triangle strip format wastes a potential factor of two in space. The geometry shown in Figure 2, for example, can be represented by one triangle strip, but many interior vertices will appear twice in the strip.

In Figure 2, a generalized triangle strip may be defined as follows, where the R denotes restart, O denotes replace oldest, M denotes replace middle, and a trailing letter p denotes push into mesh buffer. The number following a capital letter is a vertex number, and a negative number is the mesh buffer reference, in which -1 denotes the most recent pushed vertex.

R6, O1, O7, O2, O3, M4, M8, O5, O9, O10, M11, M17, M16, M9, O15, O8, O7, M14, O13, M6, O12, M18, M19, M20, M14, O21, O15, O22, O16, O23, O17, O24, M30, M29, M28, M22, O21, M20, M27, O26, M19, O25, O18 Using the same nomenclature, a generalized triangle mesh may be defined as follows:

R6p, O1, O7p, O2, O3, M4, M8p, O5, O9p, O10, M11, M17p, M16p, M-3, O15p, O-5, O6, M14p, O13p, M9, O12, M18p, M19p, M20p, M-5, O21p, O-7, O22p, O-9, O23, O-10, O-7, M30, M29, M28, M-1, O-2, M-3, M27, O26, M-4, O25, O-5

It is to be noted that a vertex reference advantageously can be considerably more compact (e.g., be represented by fewer bits) than a full vertex specification.

Three-dimensional geometry compression explicitly pushes old vertices (e.g., vertices with a trailing letter "p" above) into a queue associated with mesh buffer memory 80 (see Figure 1). These old vertices will later be explicitly referenced when the old vertex is desired again. This approach provides a fine control that supports irregular meshes of nearly any shape. Buffer memory 80 has finite length, and in practice a maximum fixed queue length of 16 is used, which requires a 4-bit index. With respect to the compression of three-dimensional graphics, the term "mesh buffer" shall refer to this queue, and the expression "generalized triangle mesh" will refer to a combination of generalized triangle strips and mesh buffer references.

The fixed size of mesh buffer 80 requires all tessellators/re-strippers for compressed geometry to break-up any runs longer than sixteen unique references. However, as geometry compression typically will not be programmed directly at the user level but rather by sophisticated tessellators/reformatters, a non-onerous restriction. Sixteen old vertices can in fact permit avoiding re-specification of up to about 94% of the redundant geometry.

Figure 2 also is an example of a general mesh buffer representation of surface geometry. Geometry compression language supports the four vertex replacement codes of generalized triangle strips, namely: replace oldest, replace middle, restart clockwise, and restart counterclockwise. Further, the language adds an additional bit in each vertex header to indicate whether or not this vertex should be pushed into the mesh buffer. In one embodiment, the mesh buffer reference command has a 4-bit field to indicate which old vertex should be re-referenced, along with the 2-bit vertex replacement code. Mesh buffer reference commands do not contain a mesh buffer push bit; old vertices can only be recycled once.

In practice, geometry rarely is comprised purely of positional data and in general, a normal, and/or color, and/or texture map coordinate are also specified per vertex. Accordingly, entries into mesh buffer 80 contain storage for all associated per-vertex information, specifically including normal and color and/or texture map coordinate.

For maximum storage space efficiency, when a vertex is specified in the data stream, per vertex normal and/or color information preferably is directly bundled with the position information. Preferably, such bundling is controlled by two state bits: bundle normals with vertices (BNV), and bundle colors with vertices (BCV). Figure 4E depicts a command structure including bits, among others. When a vertex is pushed into the mesh buffer, these bits control if its bundled normal and/or color are pushed as well.

It should be noted that the compression technique described in the above-referenced patent application is not limited to triangles, and that vectors and dots may also be compressed. Lines, for example, are a subset of triangles, in which replacement bits are MOVE and DRAW. An output vertex is then a vertex that represents one end point of a line whose other vertex is the most recently, previously omitted vertex. For dots, the replacement bits are DRAW, and

an output vertex is the vertex.

When CPU 52 executes a mesh buffer reference command, this process is reversed. That is, the two bits specify whether a normal and/or color should be inherited, or read, from the mesh buffer storage 80, or obtained from the current normal or current color. Software 58 preferably includes explicit commands for setting these two current values.

5 An exception to this rule exists, however, when an explicit "set current normal" command is followed by a mesh buffer reference, with the BNV state bit active. In this situation, the former overrides the mesh buffer normal, to allow compact representation of hard edges in surface geometry. Analogous semantics are also defined for colors, allowing compact representation of hard edges in surface colors.

10 Two additional state bits control the interpretation of normals and colors when the stream of vertices is converted into triangles. A replicate normals over triangle (RNT) bit indicates that the normal in the final vertex that completes a triangle should be replicated over the entire triangle. A replicate colors over triangle (RCT) bit is defined analogously, as shown in the command structure of Figure 4E.

15 Compression of image xyz positions will now be described. Use of the 8-bit exponent associated with 32-bit IEEE floating-point numbers allows positions to range in size from sub-atomic particles to billions of light years. But for any given tessellated object, the exponent is actually specified just once by a current modeling matrix, and object geometry is effectively described within a given modeling space using only a 24-bit fixed-point mantissa. In many cases far fewer bits are needed for visual acceptance, and the geometry compression language preferably supports variable quantization of position data down to one bit.

20 At the other extreme, empirical visual tests as well as consideration of semiconductor hardware implementation indicate that no more than 16 bits of precision per component of position is necessary for nearly all cases.

Assume, however, that the position and scale of local modeling space per object are specified by full 32-bit or 64-bit floating-point coordinates. Using sufficient numerical care, multiple such modeling spaces may be combined together to form seamless geometry coordinate systems with much greater than 16-bit positional precision.

25 Most geometry is local. Thus, within a 16-bit (or less) modeling space for each object, the difference ( $\Delta$ ) between adjacent vertices in the generalized mesh buffer stream is likely to be less than 16 bits in significance. If desired, one may construct a histogram representing bit length of neighboring position  $\Delta$ 's in a batch of geometry, and based upon this histogram assign a variable length code to compactly represent the vertices. As will be described, preferably customized Huffman coding is used to encode for the positional  $\Delta$ 's in the geometry compression.

30 Compression of red-blue-green-alpha ("RGBA") colors will now be described. Color data are treated similarly to positions, but with a smaller maximum accuracy. Thus, RGBA color data are first quantized to 12-bit unsigned fraction components that are absolute linear reflectivity values (in which 1.0 represents 100% reflectivity). An additional parameter allows color data effectively to be quantized to any amount less than 12 bits. By way of example, colors may all be within a 5-5-5 RGB color space, as shown in Figure 4C. The optional  $\alpha$  field is controlled by a color  $\alpha$  present ("CAP") state bit shown in Figure 4E. On the final rendered image individual pixel colors are still interpolated between the quantized vertex colors, and also typically are subject to lighting.

35 In practice, the same  $\Delta$ -coding may be used for color components and for positions. The area of color data compression is where geometry compression and traditional image compression confront the most similar problems. However, many advanced image compression techniques may be avoided for geometry color compression because of the difference in focus.

40 For example, the JPEG image compression standard relies upon assumptions about viewing of the decompressed data that cannot be made for geometry compression. For example, in image compression, it is known a priori that the pixels appear in a perfectly rectangular array, and that when viewed, each pixel subtends a narrow range of visual angles. By contrast, in geometry compression, the relationship between the viewer and the rasterized geometry is unpredictable.

45 In image compression, it is known that the spatial frequency of the displayed pixels upon on the viewer's eyes is likely higher than the color acuity of the human visual system. For this reason, colors are commonly converted to YUV space so that the UV color components can be represented at a lower spatial frequency than the Y (intensity) component. Usually digital bits representing sub-sampled UV components are divided among two or more pixels. However, geometry compression cannot take advantage of this because there is no fixed display scale of the geometry relative to the viewer's eye. Further, given that compressed triangle vertices are connected to four to eight or more other vertices in the generalized triangle mesh, there is no consistent way of sharing "half" the color information across vertices.

50 Similar arguments apply for the more sophisticated transforms used in traditional image compression, such as the discrete cosine transform. These transforms assume a regular (rectangular) sampling of pixel values, and require a large amount of random access during decompression.

55 It is known in the art to use pseudo-color look-up tables, but such tables would required a fixed maximum size, and would represent a relatively expensive resource for real-time processing. While pseudo-color indices could yield slightly higher compression ratios for certain scenes, the RGB model is more generalized and considerably less expensive to implement.



In an RGB model, RGB values are represented as linear reflectance values. Theoretically, if all effects of lighting could be known a priori, one or two representation bits could be dropped if the RGB components had been represented in a nonlinear, or perceptually linear space (sometime referred to as gamma corrected space). In practice, lighting effects tend not to be predictable, and on-the-fly conversion from nonlinear light to linear light would require considerable hardware resources.

The compression of surface normals will now be described. Traditionally 96-bit normals (three 32-bit IEEE floating-point numbers) are used in calculations to determine 8-bit color intensities. Theoretically, 96 bits of information could be used to represent  $2^{96}$  different normals, spread evenly over the surface of a unit sphere. The resultant extremely high accuracy represents a normal projecting in any direction every  $2^{48}$  radians.

But for IEEE floating-point normalized normals, the exponent bits are effectively unused. Given the constraint  $N_x^2 + N_y^2 + N_z^2 = 1$ , at least one of  $N_x$ ,  $N_y$ , or  $N_z$  must be in the 0.5 to 1.0 range. During rendering, this normal will be transformed by a composite modeling orientation matrix:

$$N'_x = N_x \cdot T_{0,0} + N_y \cdot T_{0,1} + N_z \cdot T_{0,2}$$

$$N'_y = N_x \cdot T_{1,0} + N_y \cdot T_{1,1} + N_z \cdot T_{1,2}$$

$$N'_z = N_x \cdot T_{2,0} + N_y \cdot T_{2,1} + N_z \cdot T_{2,2}$$

Assuming a typical implementation in which lighting is performed in world coordinates, the view transform is not involved in the processing of normals. If the normals have been pre-normalized, then to avoid redundant re-normalization of the normals, the composite modeling transformation matrix  $T$  is typically pre-normalized to divide out any scale changes. Thus:

$$T_{0,0}^2 + T_{1,0}^2 + T_{2,0}^2 = 1, \text{ etc.}$$

During normal transformation, floating-point arithmetic hardware effectively truncates all additive arguments to the accuracy of the largest component. The result is that for a normalized normal undergoing transformation by a scale preserving modeling orientation matrix, the numerical accuracy of the transformed normal value is reduced to no more than 24-bit fixed-point accuracy in all but a few special cases.

By comparison, even 24-bit normal components would still provide higher angular accuracy than the repaired Hubble space telescope, and in practice, some systems utilize only 16-bit normal components. In empirical tests with 16-bit normal components, results from an angular density of 0.01 radians between normals (e.g., about 100,000 normals distributed over a unit sphere) are not visually distinguishable from finer representations. In rectilinear space, these normals still require high representation accuracy and in practice, 16-bit components including one sign and one guard bit represents a good design choice. This still requires 48 bits to represent a normal, but since only 100,000 specific normals are of interest, theoretically a single 17-bit index could denote any of these normals.

The use of normals as indices, and the resultant advantages provided will now be described. One method of converting an index of a normal on the unit sphere back into a  $N_x$ ,  $N_y$ ,  $N_z$  value is with a table look-up, the table being loaded into memory 70 perhaps. Although table size is potentially large, the requisite size can be substantially reduced by taking advantage of a 48-way symmetry present in the unit sphere.

More particularly, as shown by Figure 3, the unit sphere is symmetrical by sign bits in the eight quadrants by sign bits. By allowing three of the normal representation bits to be the three sign bits of the xyz components of a normal, it then is only necessary to represent one eighth of the unit sphere. Each octant of the unit sphere can be divided into six identical components by folding about the planes  $x=y$ ,  $x=z$ , and  $y=z$ . The six possible sextants are encoded with another three bits, which leaves only 1/48 of the sphere remains to be represented.

Utilizing the above-noted symmetry reduces the look-up table size by a factor of  $8 \times 6 = 48$ . Instead of storing 100,000 entries, the look-up table need store only about 2,000 entries, a size small enough to be an on-chip ROM look-up table, stored perhaps within ROM 59 (see Figure 1). Indexing into the look-up table requires 11 address bits, which when added to the previously described two 3-bit fields results in a 17-bit field to describe all three normal components.

Representing a finite set of unit normals is equivalent to positioning points on the surface of the unit sphere. Although no perfectly equal angular density distribution exists for large numbers of points, many near-optimal distributions exist. Theoretically, a distribution having the above-described type of 48-way symmetry could be used for the decompression look-up table associated with the three-dimensional geometry decompression unit 130 (see Figure 1).



However, several additional constraints mandate a different choice of encoding. First, a scalable density distribution is desired, e.g., a distribution in which setting in the look-up table more low order address bits to "0" still results in fairly even normal density on the unit sphere. Otherwise a different look-up table for every encoding density would be required. Secondly, a  $\Delta$ -encodable distribution is desired in that adjacent vertices in geometry statistically have normals that are nearby on the surface of the unit sphere. Nearby locations on the two-dimensional space of the unit-sphere surface are most succinctly encoded by a two-dimensional offset. It is desirable to have a distribution in which such a metric exists. Finally, although computational costs associated with the normal encoding process are not critically important, distributions having lower encoding costs are still preferred.

Compression according to the above-referenced patent application utilizes a distribution having a regular grid in the angular space within one sextant of the unit sphere. As such, rather than a monolithic 11-bit index, all normals within a sextant are advantageously represented with two 6-bit orthogonal angular addresses. This configuration then revises the previous bit-total to 18-bits. As was the case for positions and colors, if more quantization of normals is acceptable, these 6-bit indices can be reduced to fewer bits, and thus absolute normals can be represented using anywhere from 18 to as few as 6 bits. However, as described below, this space preferably is  $\Delta$ -encoded to further reducing the number of bits required for high quality representation of normals.

Normal encoding parameterization will now be described. Points on a unit radius sphere are parameterized using spherical coordinates by angles  $\theta$  and  $\phi$ , where  $\theta$  is the angle about the y axis and  $\phi$  is the longitudinal angle from the  $y=0$  plane. Equation (1) governs mapping between rectangular and spherical coordinates as follows:

$$x = \cos\theta \cdot \cos\phi \quad y = \sin\theta \quad z = \sin\theta \cdot \cos\phi \quad (1)$$

Points on the sphere are folded first by octant, and then by sort order of xyz into one of six sextants. All table encoding takes place in the positive octant in the region bounded by the half spaces:

$$x \geq z \quad z \geq y \quad y \geq 0$$

As shown in Figure 3, the described triangular-shaped patch runs from 0 to  $\pi/4$  radians in  $\theta$ , and from 0 to a maximum 0.615479709 radians in  $\phi$ .

Quantized angles are represented by two n-bit integers  $\hat{\theta}_n$  and  $\hat{\phi}_n$ , where n is in the range of 0 to 6. For a given n, the relationship between indices  $\theta$  and  $\phi$  is:

$$\begin{aligned} \theta(\hat{\theta}_n) &= \arcsin \tan \frac{(\phi_{\max} \cdot (n - \hat{\theta}_n))}{2^n} \\ \phi(\hat{\phi}_n) &= \frac{\phi_{\max} \cdot \hat{\phi}_n}{2^n} \end{aligned} \quad (2)$$

Equations (2) show how values of  $\hat{\theta}_n$  and  $\hat{\phi}_n$  can be converted to spherical coordinates  $\theta$  and  $\phi$ , which in turn can be converted to rectilinear normal coordinate components via equation (1).

To reverse the process, e.g. to encode a given normal N into  $\hat{\theta}_n$  and  $\hat{\phi}_n$ , one cannot simply invert equation (2). Instead, the N must be first folded into the canonical octant and sextant, resulting in N'. Then N' must be dotted with all quantized normals in the sextant. For a fixed n, the values of  $\hat{\theta}_n$  and  $\hat{\phi}_n$  that result in the largest (nearest unity) dot product define the proper encoding of N. Other, more efficient methods for finding the correct values of  $\hat{\theta}_n$  and  $\hat{\phi}_n$  exist, for example indexing through the table to set  $\phi$ , and then jumping into  $\theta$ .

At this juncture, the complete bit format of absolute normals can be given. The uppermost three bits specify the octant, the next three bits the sextant, and finally two n-bit fields specify  $\hat{\theta}_n$  and  $\hat{\phi}_n$ . The 3-bit sextant field takes on one of six values, the binary codes for which are shown in Figure 3.

Some further details are in order. The three normals at the corners of the canonical patch are multiply represented, namely 6, 8, and 12 times. By employing the two unused values of the sextant field, these normals can be uniquely encoded as 26 special normals.

This representation of normals is amenable to  $\Delta$ -encoding, at least within a sextant, although with some additional work, this can be extended to sextants that share a common edge. The  $\Delta$  code between two normals is simply the difference in  $\hat{\theta}_n$  and  $\hat{\phi}_n$ , namely  $\Delta\hat{\theta}_n$  and  $\Delta\hat{\phi}_n$ .

In the above-described patent application, compression tags are used, with a variation of a conventional Huffman algorithm. The Huffman compression algorithm takes in a set of symbols to be represented, along with frequency of occurrence statistics (e.g., histograms) of those symbols. From this, variable length, uniquely identifiable bit patterns are generated that allow these symbols to be represented with a near-minimum total number of bits, assuming that

symbols do occur at the frequencies specified.

Many compression techniques, including JPEG, create unique symbols as tags to indicate the length of a variable-length data-field that follows. This data field is typically a specific-length delta value. Thus, the final binary stream consists of (self-describing length) variable length tag symbols, each immediately followed by a data field whose length is associated with that unique tag symbol.

In the referenced patent application, binary format for geometry compression uses this technique to represent position, normal, and color data fields. For geometry compression, these <tag, data> fields are immediately preceded by a more conventional computer instruction set op-code field. These fields, along with potential additional operand bits, will be referred to as geometry instructions (see Figures 4A-4K).

Traditionally, each value to be compressed is assigned its own associated label, e.g. an xyz  $\Delta$  position would be represented by three tag-value pairs. But since the  $\Delta$ xyz values are not uncorrelated, a denser, simpler representation can be attained. In general, the xyz  $\Delta$ 's statistically point equally in all directions in space. Thus, if  $n$  is the number of bits needed to represent the largest of the  $\Delta$ 's, then statistically the other two  $\Delta$  values require an average of  $n-1.4$  bits for their representation. In practice, a single field-length tag may be used to indicate the bit length of  $\Delta x$ ,  $\Delta y$ , and  $\Delta z$ .

Unfortunately, using this approach prevents taking advantage of another Huffman technique to save somewhat less than one more bit per component. However, the implemented embodiment outweighs this disadvantage by not having to specify two additional tag fields (for  $\Delta y$  and  $\Delta z$ ). A further advantage is that using a single tag field permits a hardware decompression engine to decompress all three fields in parallel, if desired.

Similar arguments hold for  $\Delta$ 's of RGB $\alpha$  values, and accordingly a single field-length tag is used to indicate bit-length of the R, G, B and, if present,  $\alpha$ , fields.

Absolute and  $\Delta$  normals are also parameterized by a single value ( $n$ ) that can be specified by a single tag. To facilitate high-speed, low-cost hardware implementations, the length of the Huffman tag field may be limited to six bits, a relatively small value. A 64-entry tag look-up table allows decoding of tags in one clock cycle. One table exists for positions, another table exists for normals, and yet another table exists for colors (and optionally, also for texture coordinates). Each table contains the length of the tag field, the length of the data field(s), a data normalization coefficient, and an absolute/relative bit.

For reasonable hardware implementation, an additional complication must be addressed. As described below, all instruction are broken-up into an eight-bit header, and a variable length body, sufficient information being present in the header to determine the body length. But the header of one instruction must be placed in the data stream before the body of the previous instruction to give the hardware time to process the header information. For example, the sequence ... B0 H1B1 H2B2 H3 ... has to be encoded as ... H1 B0 H2 B1 H3 B2 ... .

The geometry compression instruction set disclosed in the above-referenced patent application will now be described with respect to Figures 4A-4K. Figure 4A depicts a vertex command that specifies a Huffman compressed  $\Delta$ -encoded position, as well as possibly a normal and/or color, depending on bundling bits (BNV and BCV). Two additional bits specify a vertex replacement code (REP), and another bit controls mesh buffer pushing of this vertex (MBP).

As shown in Figure 4B, a normal command specifies a new current normal and the color command shown in Figure 4C depicts a new current color. The normal command and color command each use Huffman encoding of  $\Delta$  values.

The mesh buffer reference command structure is shown in Figure 4D. The mesh buffer reference command allows any of the sixteen most recently pushed vertices (and associated normals and/or colors) to be referenced as the next vertex. As further shown in Figure 4D, A 2-bit vertex replacement ("REP") code is also specified.

Figure 4E depicts the set state instruction that updates the five state bits: RNT, RCT, BNV, BCV, and CAP.

Figure 4F depicts a set table command, which is used to set entries to the entry value specified in one of the three Huffman decoding tables (Position, Normal, or Color).

Figure 4G depicts a passthrough command that allows additional graphics state not controlled directly by geometry compression to be updated in-line.

Figure 4H depicts a variable length no-op ("VNOP") command that allows fields within the bit stream to be aligned to 32-bit word boundaries. This permits aligned fields to be efficiently patched at run-time by the general CPU 52.

Figures 4I, 4J-1 and 4J-2 and 4K respectively depict tag and  $\Delta$ -position data structure, tag and  $\Delta$ -normal data structure, and tag and  $\Delta$ -color data structure. In Figures 4I and 4K, either absolute values of  $x$ ,  $y$ ,  $z$  are used, or delta values of  $x$ ,  $y$ , and  $z$  are to be used.

Of course, other instruction sets may instead be used to compress three-dimensional geometry.

The ratio of the time required for compression relative to decompression can be important. In practice, it is acceptable for off-line image compression to take up to perhaps sixty-times more time than decompression, but for real-time video conferencing, the ratio should be one.

Advantageously, geometry compression does not have this real-time requirement. Even if geometry is constructed on the fly, most geometry creating techniques, e.g., CSG, require orders of magnitude more time than needed for displaying geometry. Also, unlike continuous images found in movies, in most applications of geometry compression a compressed three-dimensional object will be displayed for many sequential frames before being discarded. Should

the three-dimensional object require animating, animation is typically done with modeling matrices. Indeed for a CD-based game, it is quite likely that an object will be decompressed billions of times by customer-users, but will have been compressed only once by the authoring company.

Like some other compression systems, geometry compression algorithms can have a compression-time vs. compression-ratio trade-off. For a given quality target level, as allowable time for compression increases, the compression ratio achieved by a geometry compression system increases. There exists a corresponding "knob" for quality of the resulting compressed three-dimensional object, and lower the quality knob, the better the compression ratio achieved.

Aesthetic and subjective judgment may be applied to geometry compression. Some three-dimensional objects will begin to appear bad when target quantization of normals and/or positions is slightly reduced, whereas other objects may be visually unchanged even with a large amount of quantization. Compression can sometimes cause visible artifacts, but in other cases may only make the object look different, not necessarily lower in quality. In one experiment by applicant, an image of an elephant actually begin to appear more realistic, with more wrinkle-like skin, as the image normals were quantized more. Once a model has been created and compressed, it can be put into a library, to be used as three-dimensional clip-art at the system level.

While many aspects of geometry compression are universal, the above-described geometry compression instruction set has been somewhat tailored to permit low-cost, high-speed hardware implementations. (It is understood that a geometry compression format designed purely for software decompression would be somewhat different.). The described geometry compression instruction set is especially amenable to hardware implementation because of the one-pass sequential processing, limited local storage requirements, tag look-up (as opposed to a conventional Hamming bit-sequential processing), and use of shifts, adds, and look-ups to accomplish most arithmetic steps.

Figure 5 is a flowchart outlining method steps in a geometry compression algorithm routine, described in the above-referenced patent application, with which the present decompression invention may be used. Such routine may be stored in memory 80 and executed under control of CPU 60 (see Figure 1).

At step 200, an object is represented by an explicit group of triangles to be compressed, along with quantization thresholds for positions, normals, and colors. At step 210, a topological analysis of connectivity is made, and hard edges are marked in normals and/or color, if such information is not already present.

At step 220, vertex traversal order and mesh buffer references are created, and at step 230 histograms of  $\Delta$ -positions,  $\Delta$ -normals, and  $\Delta$ -colors is created. At step 240, separate variable length Huffman tag codes are assigned for the  $\Delta$ -positions,  $\Delta$ -normals, and  $\Delta$ -colors, based upon histograms.

At step 250, a binary output stream is generated by first outputting Huffman table initialization, after which the vertices are traversed in order. Appropriate tags and  $\Delta$ 's are output for all values.

Applicant has implemented a Wavefront OBJ format compressor that supports compression of positions and normals, and creates full generalized triangle strips, but does not yet implement a full meshifying algorithm. Future embodiments will explore variable precision geometry, including fine structured updates of the compression tables. The current compressor expends time calculating geometric details already known to the tessellator, and ultimately it is hoped to generate compressed geometry directly. However, even its present optimized state, applicant's software can compress about 3,000 triangles/second in many cases.

The present invention is directed to decompressing three-dimensional compressed geometry, at the user end of Figure 1. Briefly, in general, an applicable geometry decompression technique according to the present invention may be outlined as follows:

- (1) Fetch the rest of the next instruction, and the first 8 bits of the following instruction;
- (2) Using the tag table, expand any compressed value fields to full precision;
- (3A) If values are relative, add to current value; otherwise replace;
- (3B) If mesh buffer reference, access old values;
- (3C) If other command, do housekeeping.
- (4) If normal, pass index through ROM table to obtain full values.
- (5) Output values in generalized triangle strip form to next stage.

In the preferred embodiment, a software embodiment of applicant's decompressor decompresses compressed geometry at a rate of about 10,000 triangles/second. A simplified overall block diagram of decompression according to the present invention is shown in figure 6. A hardware implementation of a decompressor according to the present invention can decompress in the range of tens of millions of triangles/second, which rate may be substantially expanded. Before describing decompression, it is helpful to examine and compare uncompressed and compressed image results, such as are shown in Figures 7A-7L, and in Table 1, below. Figures 7A-7G depicts the same base object, a triceratops, but with different quantization thresholds on positions and normals. Figure 7A is the original full floating-point representation, using 96-bit positions and 96-bit normals, denoted by the nomenclature P96/N96. Figure 7B and 7C depicts the effects of purely positional quantization P36/N96 and P24/N96, respectively, while Figures 7D and 7E depict only

normal quantization, P96/N18 and P96/N12. Figures 7f and 7g show combined quantization, P48/N18 and P30/N36.

Figures 7H-7L depict only quantized results for different objects, including a galleon (P30/N12), a Dodge Viper (P36/N14), two views of a '57 Chevy (P33/N13), and an insect (P39/N15).

Without zooming into the object, positional quantization much above 24-bits has essentially no significant visible effect. As the normal quantization is reduced, the positions of specular highlights on the surfaces are offset slightly. However, it is not visually apparent that such changes are reductions in quality, at least above 12 bits per normal. The quantization parameters were photographed with the objects, and otherwise even applicant could not distinguish between the original and most compressed versions of the same object.

Table 1 summarizes compression and other statistics for these objects. Column 1 notes the object in question, column 2 represents the number of  $\Delta$ 's, and column three the  $\Delta$ -strip length. The fourth column represents system overhead per vertex (overhead being everything beyond position tag/data, and normal tag/data). The "xyz quant" column denotes quantization thresholds, and the sixth column depicts the number of bits/xyz. "Bits/tri" ninth column depicts bits per triangle.

The results in Table 1 are measured actual compression data except for estimated mesh buffer results, which are shown in parenthesis. No actual mesh buffer results were present in that applicant's prototype software compressor did not yet implement a full meshifying algorithm. The estimate (in parenthesis) assumes a 46% hit ratio in the mesh buffer.

In Table 1, the right-most column shows compression ratio performance achieved over existing executable geometry formats. Although total byte count of the compressed geometry is an unambiguous number, in stating a compression ratio some assumptions must be made about the uncompressed executable representation of the object. Applicant assumed optimized generalized triangle strips, with both positions and normals represented by floating-point values to calculate "original size" data for Table 1.

To demonstrate the effect of pure 16-bit fixed point simple strip representation, Table 1 also shows byte count for the mode of OpenGL. As shown, average strip length decreased in the range of 2-3. Few if any commercial products take advantage of generalized triangle strips, and thus Table 1 considerably understates potential memory space savings.

TABLE 1

Obj. name	# $\Delta$ 's	$\Delta$ sp len.	ovrhd/vertex	xyz quant	bits/xyz	norm quant	bits/norm	bits/tri	org'l size (bytes)	comp. size (bytes)	comp. ratio
triceratops	6,039	15.9	7.5	48	30.8	18	16.8	55.9 (35.0)	179,704	42,190 (26,380)	4.3X (6.9X)
triceratops	6,039	15.9	7.5	30	17.8	12	11.0	36.0 (24.4)	179,704	27,159 (18,368)	6.7X (9.8X)
galleon	5,577	12.1	7.5	30	21.9	12	10.8	41.0 (27.2)	169,064	28,536 (18,907)	6.0X (9.0X)
Viper	58,203	23.8	7.5	36	20.1	14	10.9	37.5 (25.0)	1,698,116	272,130 (181,644)	6.3X (9.4X)
57 Chevy	31,762	12.9	7.5	33	17.3	13	10.9	35.8 (24.3)	958,160	141,830 (96,281)	6.8X (10.0X)
insect	263,783	3.0	7.5	39	22.8	15	11.0	51.5 (33.9)	9,831,528	1,696,283 (1,115,534)	5.8X (8.9X)

While certainly statistical variation exists between objects with respect to compression ratios, general trends are nonetheless noted. When compressing using the highest quality setting of the quantization knobs (P48/N18), compression ratios are typically about six. As ratios approach nearly then, most objects begin to show visible quantization artifacts.

It will be appreciated from the foregoing, that a three-dimensional geometry compression algorithm may be implemented in real-time hardware, or in software. Significantly, if three-dimensional rendering hardware contains a geometry decompression unit according to the present invention, application geometry may be stored in memory in compressed format. Further, data transmission may use the compressed format, thus improving effective bandwidth for a graphics accelerator system, including shared virtual reality display environments. The resultant compression can substantially increase the amount of geometry cacheable in main memory.

Figure 8 is a detailed block diagram of the decompressor unit 130, shown in Figure 1. As shown in Figure 8, unit 130 includes a decompression input first-in-first-out register ("FIFO") 200 whose inputs include control signals and a preferably 32-bit or 64-bit data stream, which signals and data stream preferably come from an accelerator port data FIFO ("APDF") in interface unit 120 (see Figure 1). The APDD portion of interface 120 includes a controller that signals the size of the incoming data stream to unit 130. FIFO 200 provides output to an input block state machine 220 and to an input block 210, state machine 220 and input block unit 210 communicating with each other.

Output from block 210 is coupled to a barrel shifter unit 240 and to a Huffman table set 230, the output from the Huffman look-up being coupled to state machine 220. Opcode within state machine 220 processes the values provided by the Huffman tables 230 and outputs data to the barrel shifter unit 240. State machine 220 also provides an output to data path controller 260, which outputs a preferably 12-bit wide signal to a tag decoder unit 294 and also outputs data to the barrel shifter unit 240 and to a normal processor 270, and a position/color processor 280.

Barrel shifter unit 240 outputs to the normal processor 270 and to a position/color processor 280. The outputs from processors 270 and 280 are multiplexed by output multiplexer unit 290 into a preferably 48-bit wide signal that is provided to a format converter 292.

Decompression unit 130 generates a preferably 12-bit tag that is sent to tag decoder 294 in parallel with either 32-bits or 48-bits (for normals), that are sent to the format converter 292. These data streams provide instructions that generate output to format converter 292. A preferably 32-bit read-back path is used to read-back the state of the unit.

Table 2, below, shows interface signals used to implement decompression unit 130 in the preferred embodiment:

TABLE 2

Signal Name	Signals	I/O	Description
id_data	64	I	Data inputs from APDF
id_tag	12	I	Data on inputs is valid from APDF
fd_stall	1	I	Stall signal from format converter
di_busy	1	O	Busy signal to status register
di_faf	1	O	Fifo-almost-full signal-to-input FIFO
df_data	48	O	Data output to format converter
df_tag	12	O	Tag output to tag decoder
du_context	32	O	Context output to UPA section

Table 3, below, shows output data formats provided by unit 130 in the preferred embodiment. As described herein, vertex, mesh buffer reference, and passthrough instructions generate transactions from decompression unit 130. Vertex and mesh buffer reference instructions send data to the format converter, and each generates a header indicating vertex replacement policy for the current vertex, followed by component data. Each of these instructions always generates position data and, depending upon the value of the state register, may contain color or normal data. All three of the normal components preferably are sent in parallel, whereas each position and color component is separately sent. A passthrough instruction sends preferably 32-bits of data to the collection buffer.

TABLE 3

COMPONENTS	FORMAT
Header	32.
Position	s.15

TABLE 3 (continued)

COMPONENTS	FORMAT
Color	s.15
Normal	s1.14 (x3)
Passthrough	32.

Figure 9 is a detailed block diagram of the input block 210 depicted in Figure 8. A preferably 64-bit input register 300 receives data from the APDF portion of interface 130, with 32-bits or 64-bits at a time being loaded into register 300. Register 300 outputs preferably 32-bits at a time via multiplexer 310 to a first barrel shifter 320 whose output passes through a register 330 into a merge unit 340. The 64-bit output from merge unit 340 is input to data register 350, part of whose output is returned as input to a second barrel shifter 360. The output from second barrel shifter 360 is passed through a register 370 and is also input to merge unit 340. First barrel shifter 320 aligns data to the tail of the bit-aligned data stream being recycled from data register 350 through second barrel shifter 360. The second barrel shifter 360 shifts-off the used bits from data register 350.

Figure 10 is a detailed block diagram of barrel shifter unit 240, shown in Figure 8. In overview, barrel shifter unit 240 expands the variable-length position, color, and normal index components to their fixed-point precisions. Data into unit 240 from unit 210 and/or 220 is input to a register 400 whose output is shown as defining opcode and/or data units 410, 420, 430, 440, 450, and 460, which are input to a multiplexer unit 470.

Multiplexer unit 470 input A is used for the X component of the vertex instruction, input B is used for the set normal instruction and the first component of the set color instructions, and input C is used for the remaining components of the vertex and set color instructions. Unit 240 further includes a barrel shift left register 480 coupled to receive tag\_len data and to output to register 490, whose output in turn is input to a barrel shift right register 500 that is coupled to receive data\_len data. Register 500 outputs to a mask unit 510 that is coupled to receive shift data and whose output is coupled to register 520, which outputs v\_data. The output of data block 460 is coupled to a register 530 whose output is coupled to a second register 540, which outputs pt\_data.

An appropriate table within Huffman tables 230 (see Figure 8) provides values of tag\_len, data\_len, and shift into units 480, 500 and 510, respectively. Barrel shift left unit 480 shifts the input data left by 0 to 6 bits (tag\_len), thus shifting off the Huffman tag. By contrast, barrel shift right register 500 shifts the data to the right by 0 to 16 bits (16 - data\_len), and sign extends the data, thus bringing the data to its full size. Mask unit 510 masks off the lower 'shift' bits to clamp the data to the correct quantization level.

Figure 11 depicts in greater block diagram detail the position/color processor unit 280, shown in Figure 8. Processor unit 280 generates final position or color component values. As shown in Figures 8 and 10, processor unit 280 receives a preferably 16-bit value (v\_data) from the barrel shifter unit 240, specifically mask unit 510 therein.

If the abs\_rel bit from the Huffman table 230 is set to relative, the incoming data are added by combiner unit 600 to the appropriate current stored data. The new value passes through multiplexer 610, and is stored back into the register 620, and is sent along to the output multiplexer 290, shown in Figure 8. However, if the abs\_rel bit is set to absolute, the incoming data bypasses adder 600, is latched into the register 620, and is also sent out to the output multiplexer 290.

As shown in Figure 11, the position/color processor unit 280 further includes a position/color mesh buffer 630 that is coupled to receive the input to register 620. The output from mesh buffer 630 is coupled to multiplexer gates, collectively 640, whose outputs reflect current values of x, y, z, r, g, b and  $\alpha$ . A register set, collectively shown as 650, provides these current values to the input of a multiplexer 660, whose output is coupled to the adder 600. Processor unit 280 further includes a register 670 that receives and outputs pt\_data from barrel shifter unit 240.

As shown in Figure 8, normal processor unit 270 also outputs data to the output multiplexer 290. Figure 12A depicts in detail the sub-units comprising normal processor unit 270. As seen in Figure 8 and Figure 10, the normal processor unit 270 receives an 18-bit normal index as three separate components: sextant/octant, u and v, or encoded  $\Delta u$  and  $\Delta v$  components from mask unit 510 in barrel shifter unit 240. If the value is a  $\Delta$ -value (relative), the  $\Delta u$  and  $\Delta v$  are added to the current u and v values by respective adders 710. The intermediate values are stored and are also passed on to a fold unit 800 associated with decoder-fold-rom unit 272 (see Figure 12B).

As shown in Figure 12A, the normal processor unit 270 further includes registers 712, 714, 716, 718, 720, 722, 724, 726 which hold respective octant, sextant, u and v values, curr\_oct, curr\_sext, curr\_u and curr\_v values. Also present in unit 270 are multiplexers 740, 742, 744, 746, 748, 750, 752, 754, 756, 758 and 760, its complementing units 770, 772, latch-flipflop units 780, 782, 784 for holding respective v, u, and uv information, further adders 790, 792, and a normal mesh buffer 794 coupled to receive curr\_normal input components.

With reference to Figures 12A and 12B, for an absolute reference, the u and v values are passed directly to fold



unit 800. The octant and sextant portions of the index are sent to decoder 810, within unit 272. Decoder 810 controls multiplexer 820 (which select constants), as well as multiplexers 840, 842, 844, 860, 862, 864, which reorder components, and invert signs (using 2's complement units 850, 852, 854).

Fold unit 800 uses the  $u$  and  $v$  components of the normal index, from unit 270, to calculate the address into the normal look-up table ROM 830. The octant and sextant fields, from unit 270, drive a decoder 810 that determines sign and ordering of components output from the ROM look-up table 830. Decoder 810 also handles special case normals not included in the normal ROM look-up table 830.

Figure 13 depicts interfaces to a mesh buffer, as shown in Figure 11 and/or Figure 12A. In the preferred embodiment, mesh buffer 794 is implemented as a register file and a pointer to the current location. Data is input to the mesh buffer FIFO at the position of the current location pointer. However, random access to any of the 16 locations is allowed when reading the data out of the FIFO by indexing off this pointer:  $\text{address} = (\text{curr\_loc\_ptr} - \text{index}) \bmod 16$ .

Figure 14A depicts interfaces to Huffman tables, e.g., tables 230 in Figure 8. Huffman tables are used to decode the Huffman tags preceding the compressed data. Three Huffman tables are used: one for position, for color, and for normal data, with each table preferably holding 64 entries.

Figure 14B depicts a preferred format for entry of position and color data in the Huffman tables, while Figure 14C depicts the preferred format for normal table entries. The instruction format for loading the Huffman tables in the compressed data stream is described later herein.

Several instructions generate data for the format converter 292, shown in Figure 8, and appropriate tags must be generated for this data so the format converter can correctly process the data. Table 4, below, shows tags generated for the different data components. The components that show two tags may set the launch bit, and the second tag shows the value with the launch bit set.

TABLE 4

COMPONENTS	TAG
Header	0x020
X	0x011
Y	0x012
Z	0x013/0x413
Nx/Ny/Nz	0x018/0x418
R	0x014
G	0x015
B	0x016/0x416
A	0x017/0x417
U	0x0c0/0x4c0
V	0x01c/0x41c

Input block state machine 220 (see Figure 8) includes a preferably six-bit state register that holds information about the processing state of the decompression unit. In the preferred embodiment, the following state bits are defined:

- Bit 5: tex - Texture values in place of color
- Bit 4: rnt - Replicate normal per vertex
- Bit 3: rct - Replicate color per vertex
- Bit 2: bnv - Normal bundled with vertex
- Bit 1: bcv - Color bundled with vertex
- Bit 0: cap - Color includes alpha ( $\alpha$ )

Position/Color processor unit 280 (see Figures 8 and 11) preferably includes three 16-bit registers,  $\text{curr\_x}$ ,  $\text{curr\_y}$ , and  $\text{curr\_z}$ , which contain the current position components, X, Y, and Z, and are only updated by vertex instructions.

Normal processor unit 270 (see Figures 8 and 12A) preferably includes three six-bit registers,  $\text{curr\_oct}$ ,  $\text{curr\_sext}$ ,  $\text{curr\_u}$ ,  $\text{curr\_v}$  that contain the current normal. The first register holds the 3-bit sextant and octant fields, and the remaining two registers contain the  $u$  and  $v$  coordinates for the normal. These values are written using the set normal instruction, or they are updated by the vertex instruction if the bnv bit is set in the state register.

Position/color processor 280 further preferably includes four 16-bit registers, curr\_r, curr\_g, curr\_b, curr\_a, which contain the current color components, red, green, blue and alpha ( $\alpha$ ). These components are set using the set color instruction, or they are updated by the vertex instruction if the bcv bit is set in the state register. In the preferred embodiment, alpha is valid only if the cap bit is set in the state register. The test bit is set when processing texture components, in which case only red and green are valid.

The instruction set implementing decompression according to the present invention will now be described. Figure 15A depicts the vertex instruction format, an instruction that uses variable-length Huffman encoding to represent a vertex. Position information is always present in this instruction.

(REP) The vertex replacement policy is as follows:

- 00 - Restart clockwise
- 01 - Restart counter-clockwise
- 10 - Replace middle
- 11 - Replace oldest

(M) - mesh buffer push:

- 0 - No push
- 1 - Push

With reference to Figure 15A, the position data consists of a variable-length Huffman tag (0 to 6 bits) followed by three data fields of equal length for the X, Y, and Z components, which are either  $\Delta$ -values or absolute values. The data\_len field for the entry in the position Huffman table gives the length of each of the X, Y, and Z fields, the tag\_len entry gives the length of the tag, and the abs\_rel entry tells whether the data is absolute data or is relative to the previous vertex. The shift entry from the Huffman table gives the quantization level (number of trailing zeroes) for the data.

If the bnv bit is set in the state register, a normal is included. The encoded normal has a Huffman tag followed by either two variable-length data fields for  $\Delta u$  and  $\Delta v$ , or a fixed-length field for the sextant and octant (6 bits) followed by two variable-length fields for  $u$  and  $v$ . The former encoding is for delta encodings of normals, while the latter encoding is for absolute encodings. The data\_len, tag\_len, abs\_rel, and shift fields from the normal Huffman table are used similarly as entries from the position table.

Figure 15B depicts vertex component data formats. If the bcv bit in the state register is set, color is included with the vertex. The color is encoded similar the position, using three or four fields, but how the fields are used is determined by the tag table. If tagged absolute, then x, y, z, r, g, b data is used. Absolute normals are used with sextant and octant fields. However, if the tag table indicates relative, delta normals are used, and it suffices to send latitude and longitude data (e.g.,  $\theta$  and  $\Phi$ , also referred to herein as  $u$  and  $v$ ).

With further reference to Figure 15B, a Huffman tag is followed by three equal length fields for R, G, and B. The cap bit in the state register indicates whether an additional field for  $\alpha$  is included. The data\_len, tag\_len, abs\_rel, and shift fields from the color Huffman table are used similarly as for entries from the position and normal tables.

The states of the vertex instruction set are as follows:

1. Latch next opcode; output X; shift barrel shift right unit 500 (see Figure 10) by ptag\_len + pdata\_len - pquant+2.
2. Merge; output Header.
3. Output Y; shift barrel shift right unit 500 (see Figure 10) by pdata\_len - pquant.
4. Merge
5. Output Z; shift barrel shift right unit 500 (see Figure 10) by pdata\_len - pquant.
6. Merge.

a. If (bnv)

- i. if (absolute normal), goto 7,
- ii. else goto 9. /\*relative normal\*/

b. else If (rnt), goto 21,

c. else If (bcv) goto 13,

d. else If (rct) goto 22,

e. else Merge; branch to next instruction.

7. Latch next opcode; output sextant/octant; shift barrel shift right unit 500 (see Figure 10) by ntag\_len +6.
8. Merge.
9. Output U.

- 5
  - a. If (absolute normal), shift barrel shift right unit 500 (see Figure 10) by ndata\_len - nquant.
  - b. else/\*relative normal\*/, latch next opcode; shift Bs2 by ntag\_len + ndata\_len - nquant

10. Merge.
11. Output V.
- 10 12. Merge.

- a. If (bcv), goto 13,
  - b. else If (rct), goto 22,
  - c. else Merge; branch to next instruction.

- 15 13. Latch next opcode; output R; shift barrel shift right unit 500 (see Figure 10) by ctag\_len + cdata\_len - cquant.
14. Merge
15. Output G; shift barrel shift right unit 500 (see Figure 10) by cdata\_len - cquant.
16. Merge; if (tex), branch to next instruction.
- 20 17. Output B; shift barrel shift right unit 500 (see Figure 10) by cdata\_len - cquant.
18. Merge; if (~cap) branch to next instruction.
19. Output A; shift barrel shift right unit 500 (see Figure 10) by cdata\_len - cquant.
20. Merge; branch to next instruction.
21. Output curr\_normal.

- 25
  - a. If (bcv), goto 13,
  - b. else If (rct), goto 22,
  - c. else Merge; branch to next instruction.

- 30 22. Output curr\_r.
23. Output curr\_g. If (tex), Merge; branch to next instruction
24. Output curr\_b. If (~cap), Merge; branch to next instruction.
25. Output curr\_a. Merge branch to next instruction.

35 Figure 15C depicts the format for the set normal instruction. The set normal instruction sets the value of the current normal registers. The normal data is encoded similarly as is normal data in the vertex instruction, described herein. The states of the set normal instruction are as follows:

If (absolute normal)

- 40 1. Latch next opcode; output sextant/octant; shift barrel shift right unit 500 (see Figure 10) by ntag\_len +8.
2. Merge.
3. Output U; shift barrel shift right unit 500 (see Figure 10) by ndata\_len - nquant.
4. Merge.
5. Output V; shift barrel shift right unit 500 (see Figure 10) by ndata\_len + nquant.
- 45 6. Merge; branch to next instruction.

else/\*relative normal\*/

- 50 1. Latch next opcode; output dU; shift barrel shift right unit 500 (see Figure 10) by n\_tag\_len + ndata\_len - nquant.
2. Merge.
3. Output dV; shift barrel shift right unit 500 (see Figure 10) by ndata\_len - nquant.
4. Merge; branch to next instruction.

55 Figure 15D depicts the set color instruction, an instruction that sets the value of the current color registers. Encoding of the color data is similar to encoding of the color data in the vertex instruction. The states of the set color instruction are as follows:

1. Latch next opcode; output R; shift barrel shift right unit 500 (see Figure 10) by ctag\_len + cdata\_len - cquant +2.

2. Merge.
3. Output G; shift barrel shift right unit 500 (see Figure 10) by cdata\_len - cquant.
4. Merge. If (tex), branch to next instruction.
5. Output B; shift barrel shift right unit 500 (see Figure 10) by cdata\_len - cquant.
6. Merge. If (~cap) branch to next instruction.
7. Output A; shift barrel shift right unit 500 (see Figure 10) by cdata\_len - cquant.
8. Merge; branch to next instruction.

Figure 15E is the preferred format for the mesh buffer reference instruction. This instruction causes data from an entry in the mesh buffer to be sent out to the format converter as the next vertex. With reference to Figure 15E, the index indicates the entry from the mesh buffer to send. The newest entry in the mesh buffer has index 0, and the oldest has index 15. REP, the above-described replacement policy for the vertex instruction, is the same as used for the mesh buffer reference instruction. The states for the mesh buffer reference instruction are as follows:

1. Latch next opcode; output Header; shift barrel shift right unit 500 (see Figure 10) by 9.
2. Output X from mesh buffer.
3. Output Y from mesh buffer.
4. Output Z from mesh buffer.
- a. If (bnv or rnt) goto 5,
- b. else If (bcv or rct) goto 6,
- c. else Merge; branch to next instruction.
5. If (bnv), output Normal from mesh buffer, else if (rnt) output curr\_normal.
- a. If (bnv or rct) goto 6,
- b. else Merge; branch to next instruction.
6. If (bcv), output R from mesh buffer, else if (rct) output curr\_r.
7. If (bcv), output G from mesh buffer, else if (rct) output curr\_g. If (tex), Merge; branch to next instruction.
8. If (bcv), output B from mesh buffer, else if (rct) output curr\_b. If (~cap), Merge; branch to next instruction.
9. If (bcv), output A from mesh buffer, else if (rct) output curr\_a. Merge; branch to next instruction.

Figure 15F depicts the set state instruction, which sets the bits the decompression unit state register. The states for the set state instruction are as follows:

1. Latch next opcode; shift barrel shifter 2 by 11 bits.
2. Merge; branch instruction

Figure 15G depicts the set table instruction, which sets Huffman table entries. The table selection is as follows:

- 00 - Position table
- 01 - Color table
- 10 - Normal table
- 11 - Undefined

The tag length is derived from the address. The nine bits in the entry field correspond to the absolute/relative bit, data length, and shift amount fields of the Huffman table entries. (The preferred format of the Huffman table entries has been described earlier herein.) The states of the set table instruction are as follows:

1. Latch next opcode; send address and entry to Huffman tables; shift barrel shift right unit 500 (see Figure 10) by 23.
2. Merge; branch to next instruction.

Table 5 shows the preferred Huffman Table Fill Codes.

TABLE 5

Address	Entries Filled	Tag Length	Fill Range
0ttttt	1	6	ttttt
10tttt	2	5	tttt0 - tttt1
110ttt	4	4	tttt00 - tttt11
1110tt	8	3	ttt000 - ttt111
11110t	16	2	tt0000 - tt1111
111110t	32	1	t00000 - t11111
1111110	64	0	Entire table

Figure 15H depicts the passthrough instruction, which allows passthrough data to be encoded in the compressed data stream. The length of the instruction preferably is 64-bits. Aligning successive passthrough instructions to a 64-bit boundary allows for patching of passthrough data in the encoded stream. The states for the passthrough instruction are as follows:

1. Latch next opcode; read address, shift barrel shift right unit 500 (see Figure 10) by 32 bits.
2. Merge.
3. Output data, shift barrel shift right unit 500 (see Figure 10) by 32 bits.
4. Merge; branch to next instruction.

Figure 15I depicts the variable-length NOP (\*VNOP) instruction, which encodes a variable number of 0 bits in the data stream. The five-bit count shown in Figure 15I designates the number of 0 bits that follow. This instruction is implicitly used for the start of the data stream. This instruction may also be used to pad the data stream to 32-bit or 64-bit boundaries, or encoding regions, for later patching. The states for this instruction are:

1. Latch next opcode; read count; barrel shift right unit 500 (see Figure 10) by 13 bits;
2. Merge.
3. Barrel shift right unit reads "count" positions;
4. Merge; branch to next instruction.

Figure 15J shows the skip 8 instruction, whose states are:

1. Latch next opcode; shift barrel shift right unit 500 (see Figure 10) by 16 bits;
2. Merge; branch to next instruction.

It will be appreciated that it may be advantageous to reduce bandwidth requirements between devices by not decompressing a data stream at a single point in a decompression system. An embodiment of the present invention can provide parallel decompression of a data stream by providing an additional command advising the arrival of a given number of data words that may be processed in parallel.

An embodiment of the present invention can recognize the presence of such parallel opportunities by the presence of mark bits, and cause the stated number of data words to be shuttled to other processors within the system, for parallel decompression. Further, it is then permissible to jump ahead the given number of words in the data stream to arrive at the next data that is not eligible for parallel processing.

An embodiment of the present invention can also provide morphing capability to eliminate any abrupt perception gap in viewing a decompressed three-dimensional object. Within the decompressed data stream, it is possible to specify vertices as linear or other interpolations of vertices that are actually present or have previously been decompressed. Assume, for example, that the three-dimensional object is a tiger. At a far distance, no teeth are present in the tiger's mouth, yet at near distance teeth are present. The present invention can provide a seamless transition such that as distance to the tiger shrinks, the teeth grow, with no sudden change seen between a toothless tiger and a toothed tiger.

Modifications and variations may be made to the disclosed embodiments without departing from the scope of the invention.

## Claims

1. A method for decompressing a data stream of variable length fields, said data stream including op code bits and/or tag bits, which guide decompression of said fields into uncompressed surface characteristics of three-dimensional objects, the method comprising the following steps:

(a) detecting from said data stream a field type, wherein data enabling such detecting precedes but not necessarily immediately precedes sequentially within said data stream;  
 (b) detecting from said data stream a field length and/or a field normalization, wherein data enabling such detecting precedes but not necessarily immediately precedes sequentially within said data stream;  
 (c) using information obtained from at least one of step (a) and step (b), extracting bit-aligned fields from said data stream; and  
 (d) renormalizing as necessary extracted said bit-aligned fields to an original numerical scale; and  
 (e) outputting decompressed surface-describing data, said data optionally being in a generalized triangle strip format.

2. The method of claim 1, wherein:

said op code bits represent information pertaining to type of following set of fields, said information selected from the group consisting of (i) vertex information, (ii) position information, (iii) surface normal information, (iv) color, (v) texture map coordinates, (vi) surface material properties, and (vii) procedural information, (viii) procedural information that is a command other than any of a vertex command, a normal command and a color command; and/or

said tag bits have at least one characteristic selected from the group consisting of (a) said tag bits implicitly describe at least one field characteristics selected from the group consisting of (i) field length, (ii) normalization, (iii) whether such field is relative, (iv) whether such field is absolute, and (b) said tag bits have a varying length and include information about said varying length as well as possible field content information, said varying length being independent of said field content information.

3. The method of claim 1 or 2, wherein if at step (b) the enabling data does not immediately sequentially proceed, said method further includes:

(b-1) fetching from said data stream a current instruction and a following instruction;  
 (b-2) fetching a fixed number of first bits of said following instruction; and  
 (b-3) then fetching remaining bits for said current instruction.

4. The method of any preceding claim, further including the steps of:

(f-1) providing said tag bits to a tag look-up table whose output identifies at least one of (i) start of following data fields, (ii) length of said data fields, (iii) actual tag field length, (iv) identification whether said field is relative, (v) identification whether said field is absolute, and (vi) a renormalization parameter for said data fields including numerical shifting offset information to facilitate renormalization;

(f-2) if said tag look-up table output indicates such bit-aligned fields are relative, adding a normalized value of said field to a current value of said field; and

(f-3) if said tag look-up table output indicates such bit-aligned fields are absolute, replacing a current value of said field with a said normalized value of said field.

5. The method of claim 3 or 4, wherein:

if a next following instruction is to modify how said bit-aligned fields are interpreted or is to modify contents of said tag look-up table or to pass through unmodified data in said bit-aligned fields, then do so; and  
 if such normalized bit-aligned fields contain normal information, a tag field index is passed through said look-up table to obtain a rectilinear representation of said normal.

6. The method of any preceding claim, wherein said bit-aligned fields contain information relating to at least one characteristic selected from the group consisting of (i) x,y,z coordinate positions on said surface of said object, (ii) red, green, blue color information for said object, (iii) red, green, blue and alpha color information for said object, (iv) differences between sequential red, green, blue color information for said object, and (v) differences between

sequential red, green, blue and alpha color information for said object.

7. The method of any preceding claim, wherein said object has at least one characteristic selected from the group consisting of (i) characteristics of position of said object are represented as rectilinear x, y, z coordinate, (ii) characteristics of color of said object are represented as including at least red, green, blue, and alpha components, said color being selected from the group consisting of at least (a) emissive color, (b) ambient color, (c) diffuse color, and (d) specular color, (iii) characteristics of normals of said object are represented by rectilinear  $N_x$ ,  $N_y$ , and  $N_z$  vector coefficients, (iv) characteristics of normals of said object are represented by a normal table index, said surface characteristics of said object include at least one characteristic selected from the group consisting of (a) position, (b) normals, (c) colors, (d) texture map coordinates, and (e) material surface properties.
8. The method of claim 7, wherein each of said normals is represented by a single index into a table of predefined normalized normals, and wherein said index has at least one characteristic selected from the group consisting of (i) said index includes bits of index representation comprising rectilinear representation sign bits of a normal to be represented, (ii) said index includes at least one bit specifying possible foldings of an absolute magnitude of a rectilinear representation of the normal to be represented by at least one folding plane  $x=y$ ,  $x=z$ ,  $y=z$  such that a folded normal is ensured to be on a known side of said folding plane, (iii) said index includes bits representing longitude information and latitude information of said normal, (iv) said index is defined to include bits of octant information, bits of sextant information, and additional fields of information about longitude and latitude of the folded normal, (v) said index includes bits representing longitude information and latitude information of said normal, and further includes bits of octant information, bits of sextant information, and additional fields of information about longitude and latitude of the folded normal, (vi) said index is defined to include three-bits of octant information, three-bits of sextant information, and two additional fields of information about longitude and latitude of the folded normal, and (vii) said index includes bits representing longitude information and latitude information of said normal, and further includes three bits of octant information, three-bits of sextant information, and two additional fields of information about longitude and latitude of the folded normal.
9. The method of any preceding claim, wherein said data stream to be decompressed results from said objects having been represented as three-dimensional triangular data defining a plurality of triangles having vertices, which data were converted into linear format defining a generalized triangular mesh, individual ones of said surface characteristics having been quantized, and variable length compression having been applied to a difference between sequential ones said surface characteristics.
10. The method of any preceding claim, further including the following steps:
  - if a current bit-aligned field is to be placed in a mesh buffer, placing a copy of said current bit-aligned field in said mesh buffer; and
  - if said bit-aligned fields contain a mesh buffer reference, accessing old normalized values for said bit-aligned fields.
11. A system, useable with a computer system that includes a central processing unit (CPU) and a memory coupled to said CPU, for decompressing a data stream whose fields can vary in length, said data stream including op code bits and/or tag bits, said data stream including compressed data representing three-dimensional objects whose surface defines surface characteristics, the system including:
  - a detector unit programmed to receive said data stream and to detect from a preceding, but not necessarily immediately sequentially preceding, field in said data stream field information of an immediately following field, said field information including a field type and a field length and/or field normalization; and
  - an extractor unit programmed to receive at least some data detected by said detector unit and to extract bit-aligned fields from said data stream.
12. The system of claim 11, further including a look-up table, coupled to receive from said detector unit at least a portion of said data stream, providing at least one piece of information selected from the group consisting of (i) tag field start position, (ii) actual tag field length, (iii) identification as to whether said field is relative, (iv) identification as to whether said field is absolute, and (v) numerical shifting offset information;
  - wherein bit-line fields contained in said data stream are aligned to individual bits.
13. The system of claim 11 or 12, further including at least one unit selected from the group consisting of (i) a multi-



stage barrel shifting unit that includes a first stage that aligns new incoming data contained in said data stream, and a second stage that aligns remaining old data contained in said data stream with said new incoming data, (ii) an interface circuit unit coupled to receive said data stream, (iii) a three-dimensional graphics decompressor coupled to an output of said interface circuit and decompressing said data stream, and (iv) a three-dimensional graphics rendering unit, coupled to receive and output a decompressed data stream;

wherein at least two of said barrel shifting unit, said interface circuit, said three-dimensional graphics decompressor, and said three-dimensional graphics rendering unit are fabricated on one integrated circuit.

14. The system of any of claims 11 to 13, wherein:

said op code bits represent information pertaining to type of following set of fields, said information selected from the group consisting of (i) vertex information, (ii) position information, (iii) surface normal information, (iv) color, (v) texture map coordinates, (vi) surface material properties, and (vii) procedural information, (viii) procedural information that is a command other than any of a vertex command, a normal command and a color command; and/or

said tag bits have at least one characteristic selected from the group consisting of (a) said tag bits implicitly describe at least one field characteristics selected from the group consisting of (i) field length, (ii) normalization, (iii) whether such field is relative, (iv) whether such field is absolute, and (b) said tag bits have a varying length and include information about said varying length as well as possible field content information, said varying length being independent of said field content information;

wherein at least some of said op code bits and/or said tag bits guide decompression of said fields into uncompressed surface characteristics of three-dimensional objects.

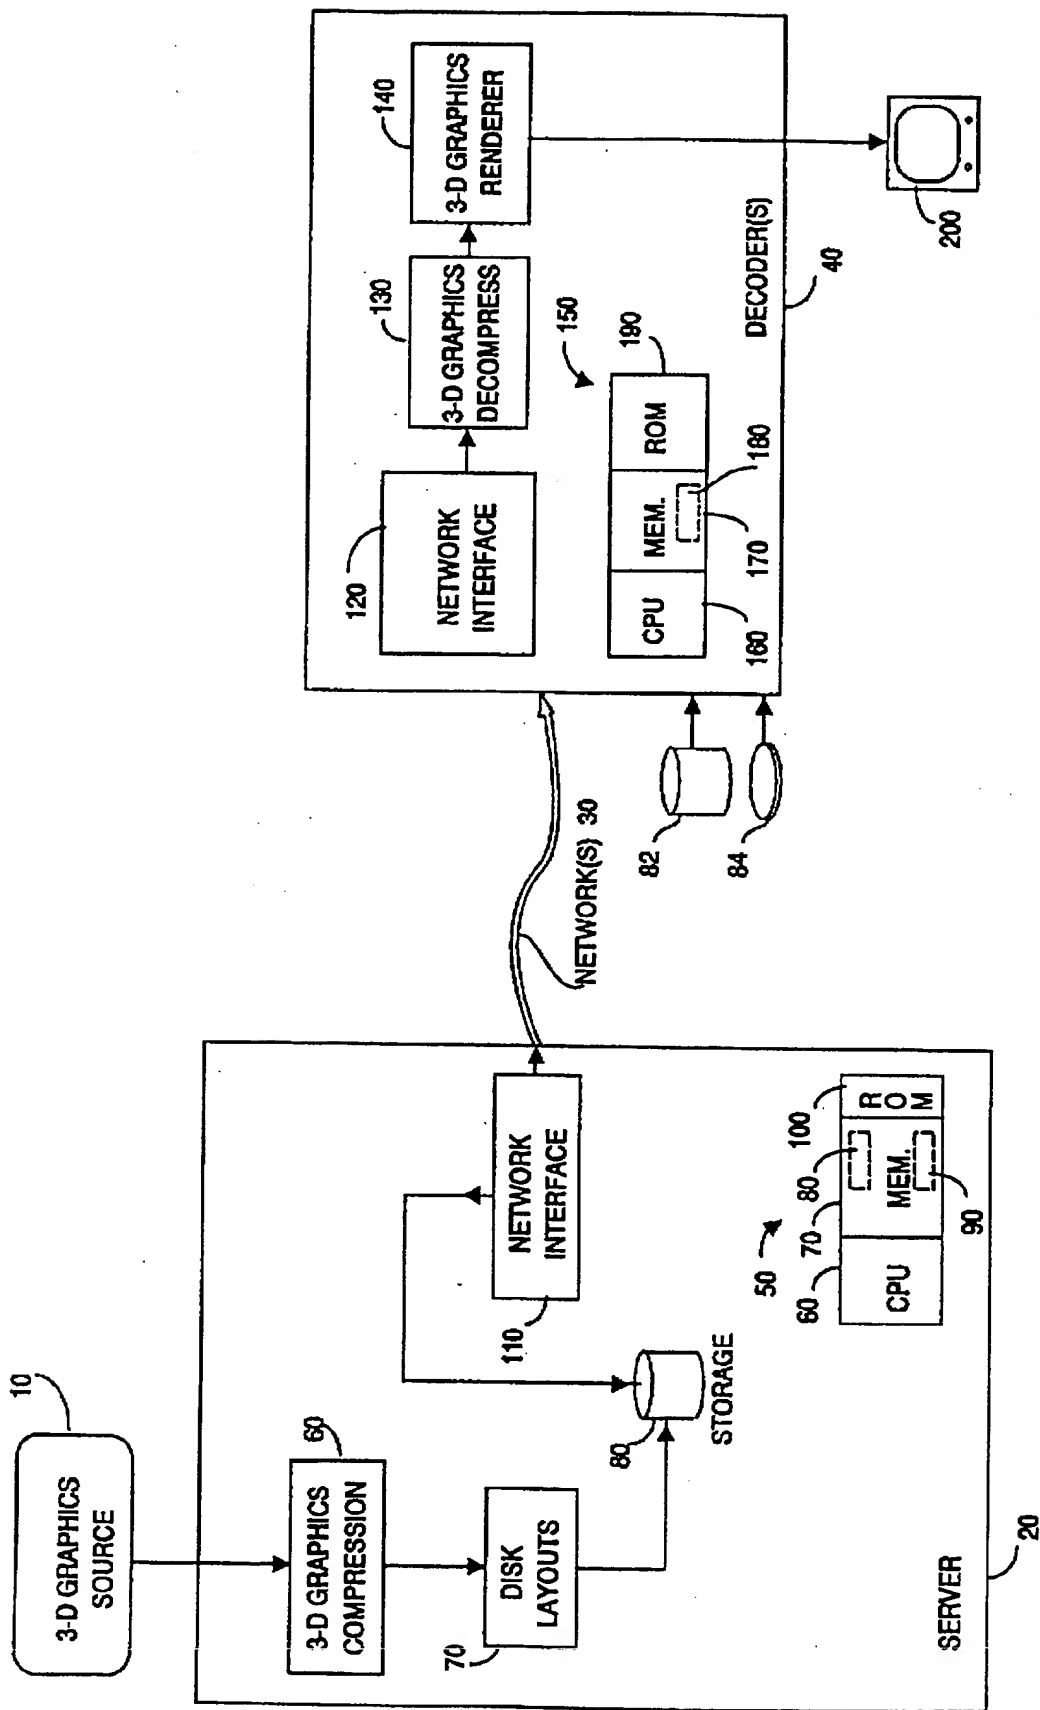
15. The system of any of claims 11 to 14, wherein:

said objects have said surface characteristics including at least one characteristic selected from the group consisting of (i) position, (ii) normals, (iii) colors, (iv) texture map coordinates, (v) material surface properties of said object; and/or


said objects have at least one characteristic selected from the group consisting of (i) characteristics of position of an object are represented as rectilinear x, y, z coordinate, (ii) characteristics of color of a said object are represented as consisting of at least red, green, blue, and alpha components, said color being selected from the group consisting of at least (a) emissive color, (b) ambient color, (c) diffuse color, and (d) specular color, (iii) characteristics of normals of a said object are represented by rectilinear  $N_x$ ,  $N_y$ , and  $N_z$  vector coefficients, (iv) characteristics of normals of a said object are represented by an index into a look-up table of said normals, (v) position, (vi) normals, (vii) colors, (viii) texture map coordinates, and (ix) material surface properties.

16. The system of claim 15, wherein said look-up table of normals contains predefined normalized normals, and wherein each of said normals is represented by a single index into said look-up table of predefined normalized normals, said index having at least one characteristic selected from the group consisting of (i) said index includes bits of index representation comprising rectilinear representation sign bits of a normal to be represented, (ii) said index includes at least one bit specifying possible foldings of an absolute magnitude of a rectilinear representation of the normal to be represented by at least one folding plane  $x=y$ ,  $x=z$ ,  $y=z$  such that a folded normal is ensured to be on a known side of said folding plane, (iii) said index includes bits representing longitude information and latitude information of said normal, (iv) said index is defined to include bits of octant information, bits of sextant information, and additional fields of information about longitude and latitude of the folded normal, (v) said index includes bits representing longitude information and latitude information of said normal, and further includes bits of octant information, bits of sextant information, and additional fields of information about longitude and latitude of the folded normal, (vi) said index is defined to include three-bits of octant information, three-bits of sextant information, and two additional fields of information about longitude and latitude of the folded normal, and (vii) said index includes bits representing longitude information and latitude information of said normal, and further includes three bits of octant information, three-bits of sextant information, and two additional fields of information about longitude and latitude of the folded normal.


17. The method of any of claims 11 to 16, wherein said data stream to be decompressed results from said objects having been represented as three-dimensional triangular data defining a plurality of triangles having vertices, which data were converted into linear format defining a generalized triangular mesh, individual ones of said surface characteristics having been quantized, and variable length compression having been applied to a difference between sequential ones said surface characteristics.




**FIGURE 1**



**FIGURE 2**

**FIGURE 3**




**FIGURE 4A**




**FIGURE 4B**



**FIGURE 4C**



**FIGURE 4D**




**FIGURE 4E**



**FIGURE 4F**



**FIGURE 4G**



**FIGURE 4H**

TAG	$\Delta X$ OR ABS X	$\Delta Y$ OR ABS Y	$\Delta Z$ OR ABS Z
-----	---------------------	---------------------	---------------------

**FIGURE 4I**

TAG	$\Delta \hat{\theta}_1$	$\Delta \hat{\phi}_1$
-----	-------------------------	-----------------------


**FIGURE 4J-1**

TAG	SECT	OCT	$\hat{\theta}$	$\hat{\phi}$
-----	------	-----	----------------	--------------

**FIGURE 4J-2**

TAG	$\Delta R$ OR ABS R	$\Delta G$ OR ABS G	$\Delta B$ OR ABS B	$\Delta \alpha$ OR ABS $\alpha$
-----	---------------------	---------------------	---------------------	---------------------------------

**FIGURE 4K**

**FIGURE 5**






FIGURE 6



**FIGURE 7A**



**FIGURE 7B**



**FIGURE 7C**




**FIGURE 7D**



**FIGURE 7E**



**FIGURE 7F**




**FIGURE 7G**




**FIGURE 7H**



**FIGURE 7I**



**FIGURE 7J**



**FIGURE 7K**



**FIGURE 7L**





FIGURE 8

**FIGURE 9**





FIGURE 10

**FIGURE 11**

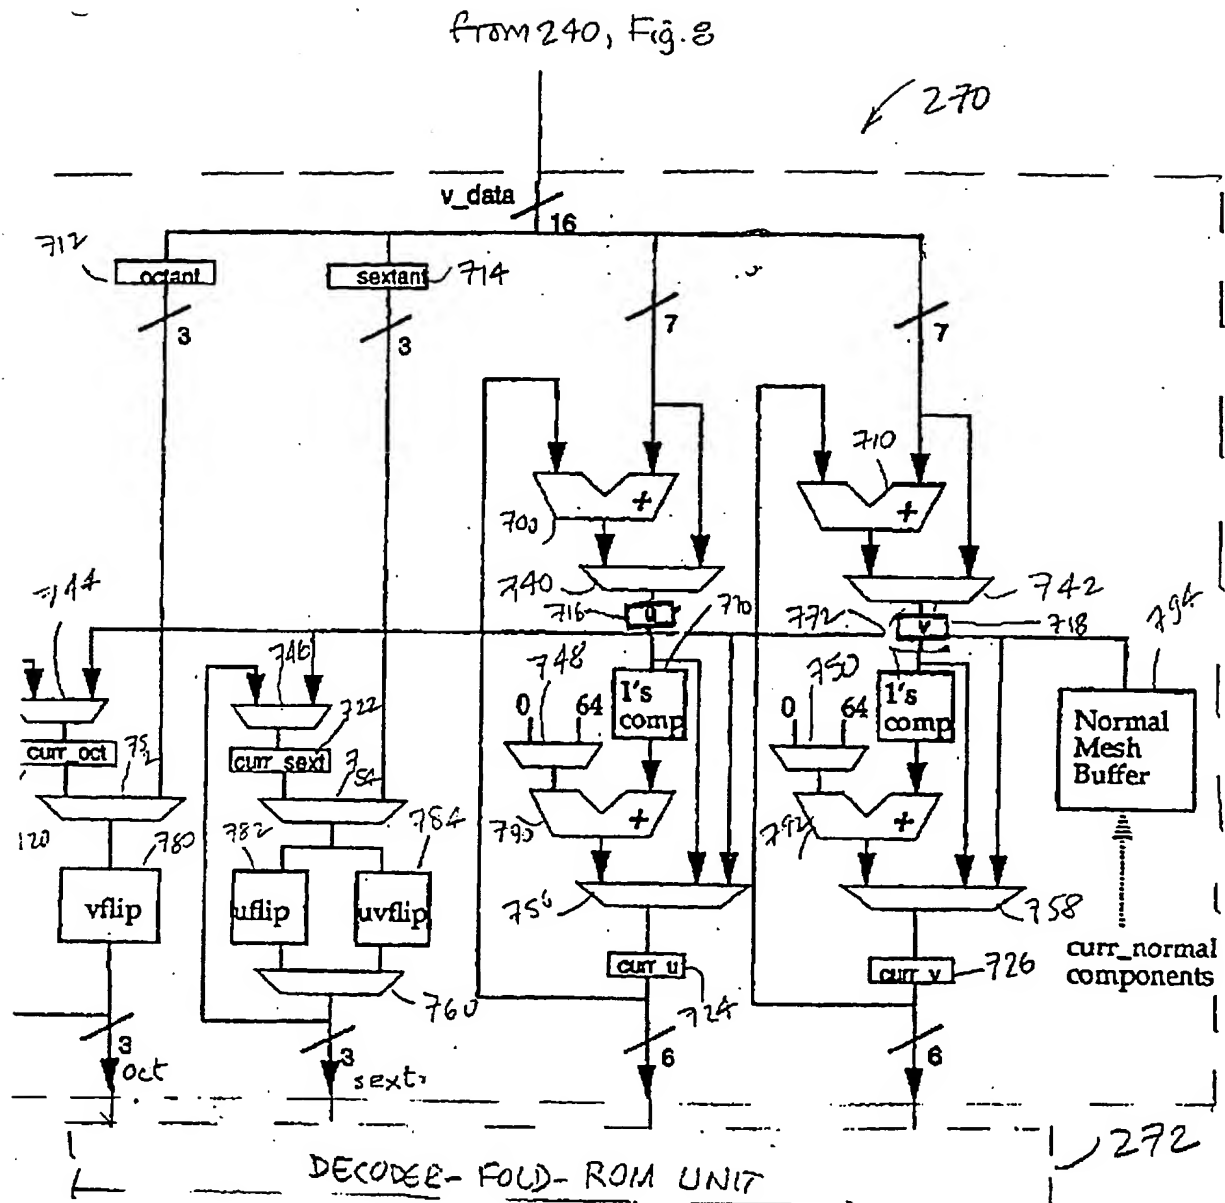

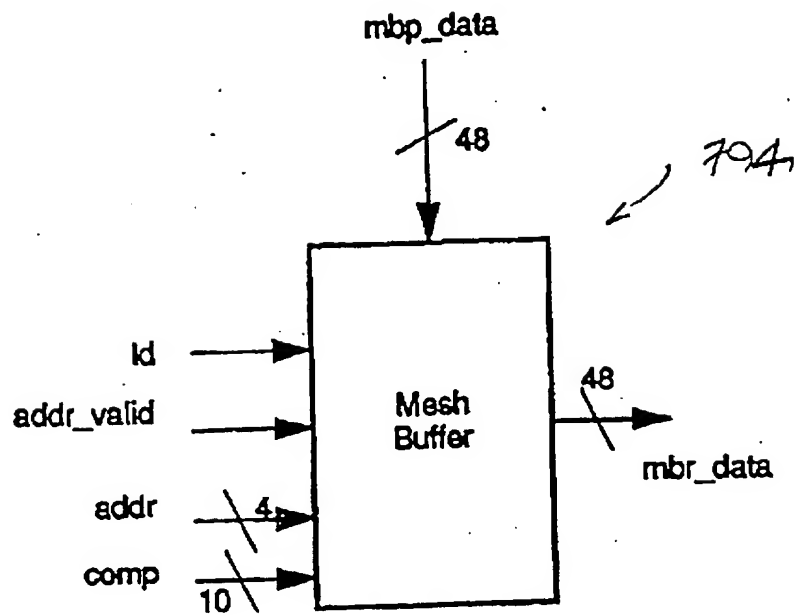




FIGURE 12A




**FIGURE 12B**


**FIGURE 13****FIGURE 14A**




**FIGURE 14B**




**FIGURE 14C**







**FIGURE 15A**





**FIGURE 15B**



**FIGURE 15C**

**FIGURE 15D****FIGURE 15E****FIGURE 15F****FIGURE 15G****FIGURE 15H**

**FIGURE 15I****FIGURE 15J**